Institute for Scientific Computing Research

Seminar Series Abstracts

(in reverse cronological order)

September 25, 2002

Open Source Python Development in SciPy and the Chaco Plotting Package

Eric Jones
Enthought, Inc.

Email: eric@enthought.com

Abstract:

Enthought, Inc., is developing the next generation of open source scientific computing tools for Python. These packages inlcude SciPy (www.scipy.org), Chaco, and Weave. SciPy provides a comprehensive toolbox of algorithms that are indespensible to scientist and engineers. It covers a wide range of areas including linear algebra, optimization, integration, special functions, signal and image processing, genetic algorithms, ODE solvers, parallel programming, and others. Plotting and 2D data visualization are provided by Chaco, a cross platform interactive plotting package. Based on DisplayPDF, Chaco is designed to work with Python, TkInter, OpenGL, and other windowing systems as well as produce high quality PDF output. Weave solves problems where pure Python just isn't fast enough. Its allow users to seamlessly integrate C/C++ code directly within Python code. This provides maximum speed with minimum fuss. Enthought is collaborating with Patrick Miller of LLNL to make this even simpler by compiling Python code to efficient C code on the fly. This talk will provide a general overview of SciPy, Chaco, and Weave as well as a discussion of future directions for the tools.

Research web page: http://www.scipy.org/

Institution web page: http://www.enthought.com/

September 17, 2002

FPGA-based Optimizations of MPI Implementations

Fred Wong
University of California, Berkeley

Email: fredwong@fredwong.com

Abstract:

Message passing libraries like MPI provide high portability across different parallel architectures, but performance suffers. Previous research has shown that performance of MPI collective operations are very sensitive to the patterns and protocols they use for communication. Therefore, researchers have put much effort into rewriting their parallel applications to avoid using high-latency collective operations like MPI reduction. Many optimized MPI implementations utilized pre-configured hardware for fast manipulation of distributed data. However, these schemes are limited fundamentally by the hardware design.

Field Programmable Gate Arrays (FPGAs) have been used traditionally as a tool for developing new architectures before mass production. Today, many systems are pre-configured with FPGA units for the sole propose of fast vector computations. This presentation gives an overview of implementing an MPI library that utilizes FPGAs versatility for collective operations with data redistribution and manipulation, as well as process synchronization. The target platform is a cluster of SMPs. Each SMP consists of an FPGA-based memory bridge connecting the host with external memory modules that can be shared with other SMP nodes. Moreover, external FPGA-based processing boards are connected to the memory modules for direct access to the shared memory.

Speaker's web page: http://www.cs.berkeley.edu/~fredwong/

Institution web page: http://www.berkeley.edu/

The Systems Biology Workbench (SBW) and the Systems Biology Markup Language (SBML)

Andrew Finney
California Institute of Technology

Email: afinney@cds.caltech.edu

Abstract:

The tremendous amounts of data and research now emerging from molecular biotechnology have fueled an explosion in the development of software tools. Regrettably, developers often end up duplicating each other's efforts when writing different packages. In an effort to make it more attractive for developers to share rather than reimplement resources, we have implemented the Systems Biology Workbench (SBW), a free, open-source, application integration environment.

SBW uses a portable architecture that enables applications (potentially running on separate computers via SSH) to learn about and communicate with each other. SBW's communications facilities allow heterogeneous packages to be connected together using a remote procedure call mechanism; this mechanism uses a simple message-passing network protocol. The interfaces to the system are encapsulated in client libraries for different programming languages (currently C, C++, Delphi, Java, Perl, and Python). SBW is portable to Windows, Linux, and MacOS X.

In this seminar, I will discuss the SBW framework and the facilities it provides for software developers. I will demonstrate some of the many SBW-enabled modules now available for users. These include a high level parser and validator for SBML; a translator from SBML to MATLAB Simulink and ODE forms; a stochastic simulator; Jarnac, an ODE-based simulator; and JDesigner, a visual tool for drawing biochemical network models. With the release of SBW version 1.0 in June 2002, development has shifted toward more module development. For example, we are developing simulation tools with support for two- and three-dimensional spatial models.

Molecular biotechnology now makes it possible to build elaborate systems models, but the systems biology community needs agreed-upon information standards if models are to be shared, evaluated and developed cooperatively. I will discuss the Systems Biology Markup Language (SBML), an open XMLbased format developed to facilitate the exchange of models of biochemical reaction networks between software packages. SBML is being developed in collaboration with the groups responsible for BioSpice, Cellerator, CellML, DBsolve, E-CELL, Gepasi, Jarnac, ProMoT/DIVA, Virtual Cell, and StochSim. Specifically, I will discuss SBML Level 2, recently developed to provide a foundation for future development of SBML. Level 2 is being developed through the close collaboration between the SBML developers forum and the DARPA BioSPICE consortium.

The SBW version 1.0 package and documentation on SBW and SBML are now available from our project web site at http://www.sbw-sbml.org/.

Research web page: http://www.sbw-sbml.org/index.html

Institution web page: http://www.caltech.edu/

September 9, 2002

Recent Advances in the TAU Performance System

Allen Malony
University of Oregon

Email: malony@cs.uoregon.edu
Sameer Shende

Abstract:

Increasing complexity in parallel and distributed systems and software necessitates advances in performance technology towards robust tools and broad implementation. However, system and software complexity introduces challenges for performance instrumentation, measurement, analysis, and visualization. This talk presents recent advances in the TAU performance system in several areas:

- selective instrumentation control
- performance mapping and dynamic callpath profiling
- online performance analysis and visualization
- performance analysis and component software
- performance databases

In the area of instrumentation control, we are concerned with the removal of instrumentation in cases of high measurement overhead using rule-based analysis. Performance mapping focuses on measuring performance with respect to dynamic calling paths in which the static callgraph cannot be determined prior to execution. A prototype online performance data access, analysis, and visualization tool is reported as an example of performance interaction and steering system. We then describe our preliminary work on performance analysis of component software based on performance technology integration in component systems.

Lastly, we introduce the TAU Performance DataBase Framework (PerfDBF). Empirical performance evaluation of parallel and distributed systems often generates significant amounts of performance data and analysis results from multiple experiments as performance is being investigated and problems diagnosed. To better manage performance information, there is a strong motivation to develop performance database technology that can provide a common foundation for performance data storage and access. Such technology could offer standard solutions for how to represent the performance data, how to store them in a manageable way, how to interface with the database in a portable manner, and how to provide performance information services to a broad set of analysis tools and users. The talk reviews a prototype PerfDBF for TAU which stores parallel performance profiles in the database and allows analysis tools to perform intra-trial, inter-trial, and cross-experiment query and analysis.

Speaker's web page: http://www.cs.uoregon.edu/~malony/

Research web page: http://www.cs.uoregon.edu/research/paracomp/tau/

Institution web page: http://www.uoregon.edu/

September 6, 2002

DMS: Software Quality Enhancement via Automated Software

Ira Baxter
Semantic Designs, Inc.

Email: idbaxter@semdesigns.com

Abstract:

Remarkably, much of software engineering today is still carried out by manual methods. Significant productivity enhancements require automation, which in turn require tools that deeply understand programs. Generative programming is a class of tool technology that captures knowledge about how to generate code, enabling automation. While generative programming is normally considered to enhance forward engineering of software, the bulk of software engineering goes into software enhancement and maintenance.

This talk will show a practical, commercial "DMS" that uses program transformations on arbitrary source languages, a kind of generative programming, to carry out a surprisingly wide variety of software enhancement tasks, including installation of test probes, duplicate code detection and removal, and automated translation from one programming language to another. Such a tool could also be used to build custom program optimizers for scientific codes.

Speaker's web page: http://www.semdesigns.com/Company/People/idbaxter/index.html

Institution web page: http://www.semdesigns.com/index.html

September 4, 2002

Filling Holes in Complex Surfaces Using Volumetric Diffusion

James Davis
Stanford University

Email: jedavis@graphics.stanford.edu

Abstract:

We address the problem of building watertight 3D models from surfaces that contain holes, for example, sets of range scans that observe most but not all of a surface. We specifically address situations in which the holes are too geometrically and topologically complex to fill using triangulation algorithms. Our solution begins by constructing a signed distance function, the zero set of which defines the surface. Initially, this function is defined only in the vicinity of observed surfaces. We then apply a diffusion process to extend this function through the volume until its zero set bridges whatever holes may be present. If additional information is available, such as known-empty regions of space inferred from the lines of sight to a 3D scanner, it can be incorporated into the diffusion process. Our algorithm is simple to implement, is guaranteed to produce manifold non-interpenetrating surfaces, and is efficient to run on large datasets because computation is limited to areas near holes.

Speaker's web page: http://graphics.stanford.edu/~jedavis/

Research web page: http://graphics.stanford.edu/~jedavis/projects/index.html

Institution web page: http://www.stanford.edu/

August 30, 2002

Programmable Graphics Hardware

Ian Buck
Stanford University

Email: ianbuck@graphics.stanford.edu

Abstract:

A recent breakthrough has occurred in graphics hardware: fixed function pipelines have been replaced with programmable vertex and fragment processors. Instead of configuring the traditional OpenGL and DirectX pipeline with awkward API calls, today's hardware allows the execution of a programmer-specified assembly level program at every fragment or vertex. This flexibility opens endless possibilities for new rendering effects and realism.

In this talk, I will outline the brief history of programmable graphics hardware as well as describe in detail how to program these new architectures. Topics covered include NVIDIA vertex programs, DirectX pixel shaders, and the CG language as well as the Radeon 9700 and the NV30 architectures. I will also present material on trends and future directions in programmability.

Speaker's web page: http://graphics.stanford.edu/~ianbuck/

Institution web page: http://www.stanford.edu/

August 23, 2002

A Modular, Parallel Grid-Embedding Adaptation Scheme for General Applications

Roger Davis
University of California, Davis

Email: davisrl@ucdavis.edu

Abstract:

The push toward high-performance aerospace designs requires the development of new advanced, fluid dynamic simulation techniques that can provide rapid and accurate predictions. Many simulation procedures used in government, industry, and academia today are based on multi-block, structured-grid systems for reasons of computational efficiency and handling of complex geometries. With these fixed-grid numerical procedures, however, an increase in solution accuracy coincides with an increase in global computational grid density and, therefore, solution cost and time. This increase in solution time becomes unacceptable as the problem size and flow complexity increases. Also, if the solution evolves in such a way that high-gradient features, such as shocks and/or high-shear viscous flows, develop in regions of coarse grid spacing, the error in the solution becomes spatially non-uniform. As a result, the solution as a whole has the accuracy associated with the region of highest numerical error. Applications, such as sonic boom or dynamic stall control that fall into this category require intelligent numerical procedures that can recognize regions of high local error and enrich the local computational region in these regions to provide the desired uniform accuracy at minimal cost.

A numerical procedure has been developed that can provide the desired level of accuracy without the penalty of large solution time through the use of grid-embedding adaptation. In this numerical procedure, computational grid cells are directionally sub-divided only in the vicinity of high numerical error. As a result, the solution evolves to one of high global accuracy and low, uniform numerical error. Since computational points are only added where they are required, the solution is obtained with minimal computational cost and time. Grid embedding is performed by sub-dividing some portion of the cells in a block depending on the desired granularity of adaptation and the areas where adaptation is desired. Sub-division of cells may be performed multiple times to create a multi-layered region of grid embedding. Sub-division of the cells may also be performed "directionally" to improve the resolution of flow features or regions where the error is proportional to the grid spacing in the direction(s) of high gradients. Adaptation granularity (i.e., the size of the region to be sub-divided) can be made arbitrary and may be controlled through the use of domain decomposition of the original block(s) into mini-blocks. By performing this additional decomposition, the overhead of the grid adaptation scheme can be kept to a minimum. Solution costs are also reduced through the use of scalable distributed memory parallel computing across groups of mini-blocks along with shared memory parallel computing within the mini-blocks themselves. This strategy is consistent with new evolving computer architectures such as the "streaming" computer.

The application of this numerical procedure is demonstrated on a number of internal and external flow configurations. In addition, details of the existing numerical algorithms, computational strategy, and data processing will be described. Finally, the extension of this adaptive-grid system to create a library of API-callable "add on" routines to allow any existing multi-block structured-grid procedure to take advantage of grid-embedding adaptation is proposed and discussed.

Institution web page: http://www.ucdavis.edu/

Some Philosophical Issues in the Foundations of Statistics

Branden Fitelson
San Jose State University

Email: branden@fitelson.org

Abstract:

What are the ends of statistical (and/or scientific) inference? What are the most effective means toward these ends? Several answers to these questions (along with further questions, of course!) will be discussed. Several well-known controversies in the foundations of statistical and scientific inference will be addressed along the way.

Speaker's web page: http://fitelson.org/

Institution web page: http://www.sjsu.edu/

July 31, 2002

Structured Population Models

Fabio Milner
Purdue University

Email: milner@math.purdue.edu

Abstract:

Some basic population models will be briefly described, both structured and unstructured. Different areas of their applicability will be discussed, including demography, mathematical epidemiology, and parasitology.

Some numerical difficulties associated with their use will be discussed, including the size of the problem, as well as that of some of the numbers involved in simulations.

Speaker's web page: http://www.math.purdue.edu/~milner/

Institution web page: http://www.purdue.edu/

July 30, 2002

Numerical Methods for Elasticity: Stress-Displacement Formulation

Zhiqiang Cai
Purdue University

Email: zcai@math.purdue.edu

Abstract:

There have been many efforts during the past four decades to develop stable mixed-finite element methods for linear elasticity. This approach is based on the Hellinger–Reissner variational principle and approximate both the displacement and the stress tensor simultaneously. Mixed methods are preferable to standard displacement methods for some important practical problems, e.g., the modelling of nearly incompressible or incompressible materials, and the modelling of plastic materials where the elimination of the stress tensor is difficult. Stable stress-displacement finite elements are extremely difficult to construct. This is caused by the symmetry constraint of the stress tensor. In this talk, I will introduce a new stable stress-displacement pair that has the least degrees of freedom and that preserves exact equilibrium for constant body force and exact symmetry of the stress tensor.

The second part of the talk is on least-squares approaches. The principal attractions of the method include freedom in the choice of finite element spaces, fast multigrid solver for the resulting algebraic equations, and free practical and sharp *a posteriori* error measure for adaptive mesh refinements. The first approach is based on the stress-displacement formulation. The second one is based on the stress-displacement-pressure formulation. The introduction of the pressure variable enables us to treat incompressible materials and to use simple standard finite element spaces approximating stress, displacement, and pressure.

Speaker's web page: http://www.math.purdue.edu/~zcai/

Institution web page: http://www.purdue.edu/

July 29, 2002

SPAN: Shared-Memory Performance Analysis

Frank Mueller
North Carolina State University

Email: mueller@cs.ncsu.edu

Abstract:

For contemporary high-performance clusters of SMPs, it has been found that a number of scientific applications utilizing a mixed mode of MPI+OpenMP are performing worse than when relying on MPI only. The objective of our work is to determine the sources of inefficiencies in utilizing memory hierarchies for threaded programs vs. parallel processes and to assist the programmer in alleviating these problems.

We are developing a technique for generating partial data traces through binary instrumentation. We employ a portable method for extracting precise data traces for partial executions of arbitrary applications and a set of hierarchical structures for compactly representing these accesses. We utilize incremental cache analysis based on partial traces to detect memory bottlenecks by correlating cache conflicts to their sources and by distinguishing intra-processor conflicts vs. inter-processor coherence misses. The conflicts will be depicted with a reference to the source program, thereby guiding the programmer to hot spots of cache misses. This enables the programmer to identify opportunities for changing the data layout and for performing algorithmic transformations to improve memory performance. Initial results on benchmarks show the potential to find memory bottlenecks and the ability of guiding the programmer to the causes of problems. These results also indicate that there is a potential to automate some optimizations in response on a running application via dynamic binary rewriting.

Speaker's web page: http://moss.csc.ncsu.edu/~mueller/

Institution web page: http://www.ncsu.edu/

July 25, 2002

Solution of Non-Symmetric, Real Positive Linear Systems

Gene Golub
Stanford University

Email: golub@sccm.stanford.edu

Abstract:

The methods we discuss use a Hermitian/skew-Hermitian splitting (HSS) iteration and its inexact variant, the inexact Hermitian/skew-Hermitian splitting (IHSS) iteration, which employs inner iteration processes at each step of the outer HSS iteration. Theoretical analyses show that the HSS method converges unconditionally to the unique solution of the system of linear equations. Moreover, we derive an upper bound of the contraction factor of the HSS iteration, which is dependent solely on the spectrum of the Hermitian part. Numerical examples are presented to illustrate the effectiveness of both HSS and IHSS iterations. In addition, a model problem of three-dimensional convection-diffusion equation is used to illustrate the advantages of our methods.

Speaker's web page: http://www-sccm.stanford.edu/Faculty/Golub.html

Institution web page: http://www.stanford.edu/

July 1, 2002

Convergence and Non-convergence in Algorithms for Prediction

Vance Faber
Mapping Science, Inc.

Email: vance@mappingscience.com

Abstract:

Many popular methods for pattern recognition and prediction involve iteration. These algorithms have been tested by using them on a wide range of problems, but no mathematical proof of convergence (in any useful sense of the word) exists. In other cases, mathematical proofs of convergence do not include information on rates of convergence.

I illustrate this situation by examining three fairly well-known algorithms: (A) Lloyd's method for clustering (a variant of k-means); (B) Breiman's algorithm for approximation of data by a hinge function; and (C) a prediction algorithm of mine which is a variant of the training stage for radial basis function neural networks. Algorithm A is known to converge sublinearly, but I construct an example to show that it is numerically non-convergent (on a finite precision machine it stalls significantly short of the solution). I analyze algorithm B and show how to produce examples where not only does it fail to converge but it actually diverges from any starting guess. Algorithm C not only converges but it can be proven that the sum of the squares of the errors that it makes will be forever bounded.

Institution web page: http://www.mappingscience.com/

OSCAR: Taking Clusters Into the Mainstream

Tim Mattson
Intel Corporation

Email: timothy.g.mattson@intel.com

Abstract:

The academic and national laboratory research communities love clusters. I doubt that there is a national lab or research university in the United States that doesn't have several of them. If you look at the corporate technical computing world, however, clusters are rare. This is slowly changing, but with few exceptions, the corporate world with its expensive labor costs can't deal with clusters today.

I believe it is in all of our interests to change this situation. We created the cluster-computing craze, and it's our job to take it into the mainstream. In this talk, I will discus one long-range effort to accomplish this; the Open Cluster Group with our first "product" OSCAR. I will describe the motivation for creating OSCAR, its current status, and our plans for the future.

Research web page: http://www.openclustergroup.org/

Institution web page: http://www.intel.com/

June 21, 2002

OpenMP: a Status Report and Look to the Future

Tim Mattson
Intel Corporation

Email: timothy.g.mattson@intel.com

Abstract:

Since its introduction in 1997, OpenMP has rapidly grown to become the standard API for writing multi-threaded applications. Its rapid transformation into a de-facto standard wasn't due to brilliance by its creators. Rather, OpenMP was a standardization of accepted practice with a narrow scope restricted to SMP machines. Coupled with support from all the standard SMP vendors, OpenMP was practically guaranteed to become a successful API.

HPC architecture marches forward and so too must the APIs that programmers use. If OpenMP wants to remain relevant, it must adapt to the increasing dominance of non-SMP systems. In addition, OpenMP needs to add constructs to support a wider range of applications. While we serve regular data driven applications quite well, we need to better handle applications that are driven by traversal of linked lists and algorithms with highly variable and unpredictable data access patterns.

In this talk, I will provide an overview of where we are today with OpenMP. I will then discuss some of the innovations we are considering as we move into the future. Finally, I plan to sit back and take notes and learn as the audience tells me what they want in future releases of OpenMP.

Research web page: http://www.openmp.org/

Institution web page: http://www.intel.com/

Shape Matching and Image Segmentation

Serge Belongie
University of California, San Diego

Email: sjb@cs.ucsd.edu

Abstract:

In this talk I will review some recent work I have done with colleagues at UC Berkeley and UCSD on shape matching and image segmentation. In the first half of the talk, I will present an approach to measuring similarity between shapes based on a novel descriptor called the "shape context." In this framework, the measurement of similarity is preceded by (1) solving for correspondences between points on the two shapes, and (2) using the correspondences to estimate an aligning transform. Object recognition results are presented for silhouettes, trademarks, handwritten digits, and 3D objects.

The preceding shape matching method assumes the objects to be recognized have been pre-segmented from the background, or that segmentation is not necessary, e.g., due to a uniform background. In the second half of the talk, I will discuss recent advances we have made on the problem of image segmentation using spectral graph theoretic methods.

Speaker's web page: http://www-cse.ucsd.edu/~sjb/

Institution web page: http://www.ucsb.edu/

June 19, 2002

JMPL: Java Math Package Launcher Automating GUI Development and Maintenance For Science and Engineering

Andrew Strelzoff
University of California, Santa Barbara

Email: strelz@engineering.ucsb.edu

Abstract:

Developers of scientific applications typically lack the time and expertise necessary to produce and maintain high quality Graphic User Interfaces (GUIs). In this presentation we introduce JMPL: Java Math Package Launcher, an environment for developing and maintaining GUI front ends for scientific computing. JMPL assists the development process by automatically generating matching GUI skeletons from scientific applications. These skeletons may be further developed and customized in the JMPL editor before distribution to the community of users. The JMPL revision manager assists in the process of updating the GUI as the underlying program changes. The result is an environment for continuing publication and improvement.

In this talk we will focus on two critical areas of the JMPL project. The first is the fundamental task of discovering interface requirements from a program and the translation of these requirements into an intermediate format, XML/UML. This intermediate format is then parsed by a browser which dynamically constructs the GUI. The second area covered will be the development of an algorithm to estimate which parts of an underlying application have changed. This is the input to the revision manager which will allow it to assist in keeping the GUI up to date. Preliminary results will be presented, and the future development and direction of the project including extension to object-oriented languages will be discussed.

Speaker's web page: http://www.cs.ucsb.edu/~strelz/

Institution web page: http://www.ucsb.edu/

June 13, 2002

Toward the Multiscale Numerical Solution of Chemically Reacting Systems

Linda Petzold
University of California, Santa Barbara

Email: petzold@engineering.ucsb.edu

Abstract:

Stochastic simulation enables the numerical simulation of the time evolution of a well-stirred chemically reacting system in a way that takes proper account of the randomness that is inherent in such a system. However, the computer times required to simulate over reasonable time periods are prohibitively long if the molecular populations of any of the reactant species are large. In cellular systems, for example, this is nearly always the case. At the same time, the deterministic reaction rate equations are well-suited for modeling of systems with large molecular populations, but cannot accurately capture the evolution of species whose populations are small.

In this talk we will examine a range of models from stochastic simulation to the deterministic reaction rate equations. Since the models naturally segue to each other, there is some reason to hope that a multiscale numerical method that treats each species and reaction at the most appropriate scale could be developed. We will describe our progress on the first phase of this effort: the development of efficient and robust accelerated stochastic simulation algorithms inspired by the recently-developed tau-leaping method of Gillespie.

This is joint work with Dan Gillespie and Muruhan Rathinam.

Speaker's web page: http://www.me.ucsb.edu/dept_site/people/new_faculty_pages/petzold_page.html

Institution web page: http://www.ucsb.edu/

June 12, 2002

Multi-Implicit Methods for Advection-Diffusion-Reaction and Related Problems

Michael Minion
University of North Carolina, Chapel Hill

Email: minion@amath.unc.edu

Abstract:

A new method for the temporal integration of partial differential equations possessing multiple time scales will be presented. The method is based on a multi-implicit variation of spectral deferred corrections (MISDC) in which two or more components of the solution are treated implicitly. As in traditional operator splitting methods, when applied to Advection-Diffusion-Reaction (A-D-R) equations, MISDC methods allow both the diffusion and the reaction terms to be treated in an implicit but uncoupled manner. Furthermore, the size of the time steps used for the the diffusion and reaction terms can be chosen independently. Higher-order temporal accuracy is achieved by reducing both the splitting and truncation errors during the deferred correction iterations. Examples of MISDC method of order 3, 4, and 5, applied to model A-D-R problems will be presented. The difficulties in imposing accurate boundary conditions for time-dependent boundary conditions will also be discussed.

Speaker's web page: http://www.amath.unc.edu/Faculty/minion/

Institution web page: http://www.unc.edu/

June 11, 2002

Multigrid for Eddy Current Computation

Ralf Hiptmair
Universitaet Tuebingen

Email: hiptmair@na.uni-tuebingen.de

Abstract:

I set out from the quasi-state eddy current model in time domain. Because of stability, implicit time-stepping in required. In a finite element setting, each timestep involves the solution of a discrete variational problem for a bilinear form (•, •) + (curl •, curl •) posed over H(curl). I rely on Nedelec's H(curl)-conforming finite elements (edge elements), which yield a viable discretization for Maxwell's equations. Ultimately, we are faced with a large sparse linear system of equations to be solved in each timestep, for which we aim to construct a fast iterative solver.

A naïve multigrid approach is doomed, since the problem is hardly elliptic on the large kernel of the curl-operator. Hence, plain nodal relaxations are only effective for high-frequency error components that are sufficiently orthogonal to this kernel. Yet, the exceptional properties of edge elements provide discrete potentials in the space of piecewise polynomial Lagrangian finite element functions. Thus, additional smoothing in the kernel becomes computationally feasible and, ultimately, leads to a multigrid method, which achieves “textbook multigrid efficiency.”

Speaker's web page: http://na.uni-tuebingen.de/~hiptmair/

Institution web page: http://www.uni-tuebingen.de/

June 10, 2002

Higher Order Whitney Forms

Ralf Hiptmair
Universitaet Tuebingen

Email: hiptmair@na.uni-tuebingen.de

Abstract:

The calculus of differential forms can be used to devise a unified description of discrete differential forms of any order and polynomial degree on simplicial meshes in any spatial dimension. A general formula for suitable degrees of freedom is also available. Fundamental properties of nodal interpolation can be established easily. It turns out that higher order spaces, including variants with locally varying polynomial order, emerge from the usual Whitney-forms by local augmentation. This paves the way for an adaptive p-version approach to discrete differential forms.

Speaker's web page: http://na.uni-tuebingen.de/~hiptmair/

Institution web page: http://www.uni-tuebingen.de/

May 29, 2002

Compile-time Analysis of Loop Nest Behavior in Cache Memory

Erin Parker
University of North Carolina, Chapel Hill

Email: parker@cs.unc.edu

Abstract:

It is difficult to talk about high-speed computing without considering the bottleneck of accessing memory. For large applications, achieving good performance means taking full advantage of fast, small cache memories. Applications exhibiting good locality realize the benefits of caches; however, locality is often elusive and volatile. As an application evolves, modifications to its code and platform lead to changes in its locality. The need for efficient, analytical models of the memory behavior of programs is evident.

We provide an exact model of the behavior of loop nests executing in a memory hierarchy. Our model is flexible enough to describe whether a memory access hits or misses in the cache, whether a memory block experiences reuse while in the cache, and at what point a memory block experiences no more reuse before being evicted from the cache. We use Presburger arithmetic to express the scenarios above and a combination of tools to simplify and identify such memory accesses. We can then use this information to predict the number of cache misses incurred by loop nests, to avoid polluting the cache with memory blocks that are not reused, to choose when to evict memory blocks from cache in order to reduce bus congestion, and even to turn off cache lines holding data not reused to reduce cache leakage power.

Speaker's web page: http://www.cs.unc.edu/~parker/

Institution web page: http://www.unc.edu/

May 24, 2002

Fast and Stable Direct Solvers for Matrices from PDEs and Integral Equations

Shivkumar Chandresakaran
University of California, Santa Barbara

Email: shiv@ece.ucsb.edu

Abstract:

We will present the "sequentially semi-separable (SSS)" class of dense structured matrices. This class includes dense matrices, semi-separable matrices, the inverse of both dense and semi-separable matrices, and many others. We will present fast and stable algorithms for the solution of linear systems of equations with SSS matrices. We will also present fast and memory efficient techniques for constructing the SSS representation. Finally we will show how two spectral methods, one for two-point boundary value problems and one for exterior two-dimensional scattering problems, can be solved in linear time using the SSS representation. Experimental results on the efficiency of the fast algorithms will also be presented.

This is joint work by S.Chandrasekaran, M.Gu, P.Dewilde, T.Pals, A.-J.van der Veen.

Speaker's web page: http://www.ece.ucsb.edu/Faculty/Chandrasekaran/default.html

Institution web page: http://www.ucsb.edu/

May 23, 2002

Cache-Oblivious Priority Queue and Graph Algorithm Applications

Lars Arge
Duke University

Email: large@cs.duke.edu

Abstract:

In recent years algorithms designed to take advantage of the hierarchical memory system of modern machines have been extensively researched. Most research on such algorithms has been done in two-level (or external memory) memory models. The advantage of such models is that they are simple, while the disadvantage is that they focus efforts on one level of the multilevel hierarchy. Recently a new model that combines the simplicity of two-level models with the realism of more complicated and realistic hierarchical models was introduced. The idea in this "cache oblivious" model is to design algorithms in a two-level model but avoid using the parameters describing the two-level hierarchy in the algorithm description. This way an algorithm that is efficient in the two-level model is efficient on any level of an arbitrary multilevel memory hierarchy.

In this talk we first discuss fundamental sorting and searching results in the cache-oblivious model. Then we describe an optimal cache-oblivious priority queue data structure, supporting insertion, deletion, and deletemin operations in O((1/B) log(M/B) (N/B)) amortized memory transfers, where M and B are the memory and block transfer sizes of any two consecutive levels of a multilevel memory hierarchy. These bounds match the bounds of several previously developed external-memory (cache-aware) priority queue data structures, which all rely crucially on knowledge about M and B. Priority queues are a critical component in many of the best known external-memory graph algorithms, and the cache-oblivious priority queue can, for example, be used to develop several cache-oblivious graph algorithms. In most cases the memory-transfer bounds of these algorithms match the bounds of the best known external-memory algorithms.

Joint work with Michael Bender, Erik Demaine, Bryan Holland-Minkley and Ian Munro presented at STOC'02.

Speaker's web page: http://www.cs.duke.edu/~large/

Institution web page: http://www.duke.edu/

May 17, 2002

Incrementally Mining Frequent Patterns and Mining Bio-medical Datasets

Srinivasan Parthasarathy
Ohio State University

Email: srini@cis.ohio-state.edu

Abstract:

In the first part of the talk I will describe our work in incrementally mining frequent patterns in evolving datasets. Changes to a dataset can invalidate existing patterns or introduce new ones. Simply re-executing algorithms from scratch on a database update can result in an explosion in the computational and I/O resources required. What is needed is a way to incrementally process the data and update the information that is gleaned while being cognizant of the interactive requirements of the knowledge discovery process.

During the second part of the talk I will describe some recent work in mining biomedical datasets. A common theme across the proposed techniques here is that they are cognizant of the structure/shape-function relationship that pervades a lot of applications in these domains. I will describe techniques that are cognizant of features relating to the shape and structure of biological elements (corneal shapes, protein structures) and how these techniques can be used in conjunction with classical data mining techniques to mine such datasets.

Acknowledgements: The first part of the talk is joint work with researchers from RPI and UFMRG (Brazil). The second part of the talk is joint work with graduate students at OSU.

Speaker's web page: http://www.cis.ohio-state.edu/~srini/

Research web page: http://www.cis.ohio-state.edu/~srini/

Institution web page: http://www.ohio-state.edu/index.php/

May 10, 2002

Data Reduction Techniques using Subsampling

Lori Freitag and Raymond Loy

Argonne National Laboratory

Email: freitag@mcs.anl.gov

Abstract:

The interactive visualization and exploration of large scientific data sets is a challenging and difficult task; their size often far exceeds the performance and memory capacity of even the most powerful graphics workstations. To address this problem, we have created a technique that combines hierarchical data reduction methods with parallel computing to allow interactive exploration of large data sets while retaining full-resolution capability. Our technique is based on subsampling the original data set using a distributed octree, and we describe the software architecture of the system in some detail. We develop performance models for this and uniform subsampling strategies to investigate the tradeoffs between multiresolution and uniform grid methods in terms of performance and approximation errors.

Speaker's web page: http://www-unix.mcs.anl.gov/~freitag/

Research web page: http://www-unix.mcs.anl.gov/~freitag/

Institution web page: http://www.anl.gov/

April 23, 2002

Future Technologies Activities at LBL

Brent Gorda
Lawrence Berkeley National Laboratory

Email: bgorda@lbl.gov

Abstract:

Brent Gorda, newly appointed group lead for the Future Technologies Group at the Lawrence Berkeley National Laboratory, will talk about the activities of the group in the area of High Performance Computing.

Research web page: http://www.nersc.gov/research/ftg/

Institution web page: http://www.lbl.gov/

April 18, 2002

An Overview Of Verification and Validation for Computational Science and Engineering

Patrick Roache
Ecodynamics Research Associates, Inc.

Email: hermosa@swcp.com

Abstract:

This seminar presents an overview of Verification and Validation for Computational Science and Engineering and related areas. Topics include: semantics and definitions; Verification of Codes by the Method of Manufactured Solutions (MMS); Partitioning the Option Matrix; Verification of Calculations using the Grid Convergence Index (GCI); Single-Grid Error Estimators; validation experiments and metrics; "nearby" problems; scale of unsteadiness; scaling-up; prevalence of errors in scientific software; use of static analyzers in QA systems; and making QA your friend.

Speaker's web page: http://kumo.swcp.com/hermosa/html/pjr.html

April 12, 2002

An Immersed Interface Method for Incompressible Navier–Stokes Equations

Long Lee
University of Washington

Email: longlee@amath.washington.edu

Abstract:

Peskin's immersed boundary (IB) method is one of the most robust numerical methods for solving fluid problems with immersed interfaces. In this talk, I will present a method motivated by the immersed boundary method, which allows one to model the motion of flexible membranes or other structures immersed in viscous incompressible fluid using a fluid solver on a fixed Cartesian grid. The IB method uses a set of discrete delta functions to spread the entire singular force exerted by the immersed boundary to the nearby fluid grid points. Our method instead incorporates part of this force into jump conditions for the pressure, avoiding discrete dipole terms that adversely affect the accuracy near the immersed boundary. This has been implemented for the two-dimensional incompressible Navier–Stokes equations using a high-resolution finite-volume method for the advective terms and a projection method to enforce incompressibility. In the projection step, the correct jump in pressure is imposed in the course of solving the Poisson problem. This gives sharp resolution of the pressure across the interface and also gives better volume conservation than the traditional IB method. Comparisons between this method and the IB method are presented for several test problems. Numerical studies of the convergence and order of accuracy are included. This work is joint with Randy LeVeque.

Speaker's web page: http://amath.washington.edu/~longlee/

Institution web page: http://www.washington.edu/

April 3, 2002

Numerical Methods for Air Pollution Modeling and Simulation

Florian Potra

University of Maryland, Baltimore County

Email: potra@math.umbc.edu

Abstract:

A detailed understanding of the relationships between the emissions and the resulting distribution of primary and secondary species in the atmosphere is a requisite to designing actions for the maintenance of a healthy environment. To better predict the transport and fate of trace gases and pollutants in the atmosphere, comprehensive atmospheric transport-chemistry models have been developed. These models treat transport, chemical transformations, emissions and deposition processes in an integrated framework. They can be used to study the response of the pollutant distributions to system perturbations, and to link pollutant distributions to environmental effects. However, these comprehensive atmospheric chemistry models are computationally intensive because the governing equations are nonlinear, highly coupled, and extremely stiff. Our ability to fully utilize these models remains severely limited by today's computer technology and algorithmic development.

The talk will focus on improved numerical methods for solving the stiff differential equations associated with the chemical reactions and on new software tools for generating the computer code.

Speaker's web page: http://www.math.umbc.edu/~potra/

Institution web page: http://www.umbc.edu/

March 26, 2002

Cooperative Research at IBM's Centre for Advanced Studies (CAS)

Gabriel Silberman
IBM, TJ Watson Laboratory

Email: gabbys@us.ibm.com

Abstract:

The Centre for Advanced Studies (CAS) was established by the IBM Toronto Laboratory in 1991 to strengthen the links between academia and industry. Since then, this research group has been helping bridge the gap between industry requirements for new technology and academia.

CAS places an emphasis on solving real problems encountered by developers. These lead to long-term projects from which prototypes and publication-quality research can be delivered, all in collaboration with other members of the research community.

Speaker's web page: http://www.cas.ibm.ca/director/

Research web page: http://www.cas.ibm.ca/index.shtml

Institution web page: http://www.watson.ibm.com/

March 18, 2002

Scalable On-line Automated Performance Diagnosis

Philip Roth
University of Wisconsin–Madison

Email: pcroth@cs.wisc.edu

Abstract:

Automated search is an effective strategy for finding application performance problems. Search-based automated tools are especially important for performance diagnosis of large-scale applications. Unfortunately, existing automated performance diagnosis tools do not scale well.

In the first part of this talk I present Deep Start, a new algorithm that attacks the scalability problem of automated performance diagnosis for applications with a large number of functions. Deep Start uses stack sampling to augment a simple search-based automated performance diagnosis strategy. Our hybrid approach locates performance problems more quickly and finds problems hidden from a more straightforward search strategy. Deep Start uses stack samples collected as a by-product of normal search instrumentation to find deep starters, functions that are likely to be application bottlenecks. Deep Starters are examined early during a search to improve the likelihood of finding performance problems quickly. We implemented the Deep Start algorithm in the Performance Consultant, the automated bottleneck detection component of the Paradyn performance tool. Deep Start found half of our test applications' known bottlenecks 32% to 59% faster than the Performance Consultant's current call graph-based search strategy, and finished finding bottlenecks 10% to 61% faster than the call graph strategy. In addition to improving search time, Deep Start often found more bottlenecks than the call graph search strategy.

In the second part of the talk, I will outline our plans for improving the scalability of automated performance diagnosis tools for applications that create a large number of processes. I will discuss our strategy for distributing the Performance Consultant, including our design for a general-purpose software-based multicast/reduction network. I will also present the Sub-graph Folding Algorithm, our approach for improving the scalability of the Performance Consultant's results display.

Speaker's web page: http://www.cs.wisc.edu/~pcroth/

Research web page: http://www.paradyn.org/

Institution web page: http://www.wisc.edu/

March 8, 2002

A Hybrid Particle Level Set Method for Improved Interface Capturing

Doug Enright
Stanford University

Email: enright@stanford.edu

Abstract:

In this talk, a new numerical method for improving the area (volume in three spatial dimensions) conservation properties of the level set method when the interface is passively advected in a flow field is presented. This method uses Lagrangian marker particles to rebuild the level set in regions that are under-resolved as is often the case for flows undergoing stretching and/or tearing. The overall method maintains a smooth geometrical description of the interface and the implementation simplicity characteristic of the level set method. The method compares favorably with volume of fluid methods in the conservation of mass and purely Lagrangian schemes for interface resolution in highly vortical flows. Numerical examples in both two and three spatial dimensions are shown along with preliminary results of the use of this method in free-surface calculations.

Speaker's web page: http://www-sccm.stanford.edu/~enright/

Research web page: http://cits.stanford.edu/

Institution web page: http://www.stanford.edu/

March 5, 2002

Smarter Memory Controllers: Improving Memory System Performance from the Bottom Up

Sally McKee
University of Utah

Email: sam@cs.utah.edu

Abstract:

Microprocessor speed is increasing much faster than memory system speed: the speed-growth curves are diverging exponentials. The traditional approach to attacking the memory system bottleneck has been to build deeper and more complex cache hierarchies.Although caching may work well for parts of programs that exhibit high locality, many important commercial and scientific workloads lack the locality of reference that makes caching effective. Studies have shown that memory bus and DRAM latencies cause an 8X slowdown from peak performance to actual performance on commercial database workloads, that the efficiency of current caching techniques is generally less than 20% of an optimal cache's, and that cache sizes are up to 2,000 times larger than an optimal cache would be. The evidence is clear: no matter how hard we push it, traditional caching cannot bridge the growing processor-memory performance gap.

This talk presents research that attacks the memory problem at a different level—the memory controller. We describe the Impulse Adaptable Memory System being built at the University of Utah. This general-purpose, uniprocessor system that improves performance within the cache hierarchy and the memory back end for both regular and irregular computations. It does this in three ways: by optimizing the use of DRAM resources in the memory controller back end, by prefetching data within the memory controller and delivering it to the processor only when requested, and by remapping previously unused physical addresses within the memory controller. Extending the virtual memory hierarchy in this way allows optimizations that improve the efficiency of the system bus and the performance of the CPU's cache hierarchy and translation lookaside buffer (TLB). Impulse represents a combined hardware/software solution in that the compiler, OS, or application writer supplies the access pattern information that the memory controller exploits.

We provide an overview of Impulse functionality, and then present details of the Parallel Vector Access (PVA) mechanism that optimizes the use of DRAM resources by gathering data in parallel within the memory controller. The PVA performs cache-line fills as efficiently as a normal, serial controller, and performs strided vector accesses from three to 33 times faster. The scalable design is two-five times faster than other-gathering mechanisms with similar goals, at the cost of only a slight increase in hardware complexity.

Speaker's web page: http://www.cs.utah.edu/~sam/

Institution web page: http://www.utah.edu/

February 28, 2002

Fast Wavelet-based Evaluation of Polynomial Range-Sum Queries

Rolfe Schmidt
University of Southern California

Email: rrs@pollux.usc.edu

Abstract:

Many range aggregate queries, including COUNT, SUM, and COVARIANCE, can be efficiently derived from a fundamental class of queries: the polynomial range-sums. This talk will describe how range-sum queries can be thought of as linear operators on the data density function, and evaluated exactly in the wavelet domain. With this intuition, I will introduce ProPolyne, a novel technique to evaluate arbitrary polynomial range-sums progressively. At each step of the computation, ProPolyne makes the best possible wavelet approximation of the submitted query. When the filters are chosen appropriately, polynomial range-sum queries have very sparse wavelet representations. The result is an excellent exact-query answering technique that provides highly accurate data-independent approximate results long before query evaluation is finished. I will conclude by discussing a few of the challenges that must be addressed to fully exploit this technique, and our progress in overcoming them.

Research web page: http://dimlab.usc.edu/research.html

Institution web page: http://www.usc.edu/

February 27, 2002

Mesh-based Modeling Methods for Computational Biology Applications

Harold Trease
Pacific Northwest National Laboratory

Email: het@pnl.gov

Abstract:

Biological systems are inherently multi-dimensional because of their geometry, heterogeneity, and time-dependent behavior. One approach to capturing these aspects is to use mesh-based modeling methods as the basis for a computational simulation framework that explicitly represents complex geometry, non-uniform heterogeneous distributions of concentrations, time-dependence, and complex boundary/initial conditions. To build such simulation models there are several problems that must be solved to integrate biology, biophysics, mathematics, and computer science into a useable model.

This presentation demonstrates the approaches that we at the Pacific Northwest National Laboratory are taking toward creating systems biology models from everything from computational cell biology to complex organ models. Specifically, we will show the virtual microbial cell and virtual mammalian nose/lung models that we are constructing. This includes using various forms of digital image reconstruction and feature extraction for NMR/MRI, confocal, and ultrasound to obtain computation geometry models, from which our mesh models are generated. We also show the mathematical models that we are using to simulate the combined effects of fluid-dynamics, tissue mechanics, reaction/diffusion, and discrete particle transport for these multi-component biological systems. The simulation framework presented is part of the DOE SciDAC Terascale Simulation Tools and Technology (TSTT) Center. Both LLNL and PNNL are contributing members of this center.

Institution web page: http://www.pnl.gov/

February 21, 2002

Parallel Performance Diagnosis with PperfDB

Karen Karavanic
Portland State University
Mathew Colgrove
Email: karavan@cs.pdx.edu

Abstract:

The PPerfDB research project is developing methods for diagnosing large-scale parallel applications using data from more than one program execution. This multi-execution view enables comparison of different code versions, different communication libraries, different platforms, and even different monitoring tools. The motivating vision driving the PPerfDB project is a tool that can diagnose the performance of large-scale parallel applications automatically, without user involvement.

PPerfDB can launch new performance experiments using IBM's Dynamic Probe Class Library (DPCL). PPerfXchange, a tool component under development, allows geographically dispersed data stores to be included as sources of data for PPerfDB. PPerfXchange models each site's performance data as XML documents, based upon a global schema, allowing client applications to retrieve views on performance data using the XQuery Language.

In this talk we will describe PPerfDB and report results that demonstrate our prototype's ability to combine performance data from traditional tracing tools and dynamic instrutmention-based tools for comparative performance profiling.

Speaker's web page: http://www.cs.pdx.edu/~karavan/research.html

Institution web page: http://www.pdx.edu/

February 7, 2002

Whats New with the Wave Equation

Thomas Hagstrom
University of New Mexico

Email: hagstrom@math.unm.edu

Abstract:

In recent years there have been a number of subtantial advances in numerical techniques for solving the wave equation. These include:

(i) Efficient methods for truncating unbounded domains which can provide arbitrary accuracy.

(ii) Fast methods for evaluating the standard integral solution formulas, akin to the fast multipole method in the frequency domain.

(iii) New discretization strategies including stabilized one-sided differencing near boundaries and unconditionally stable, explicit 2-step marching methods based on new integral formulas for the solution.

I will review these new methods, discuss their applicability to more general hyperbolic initial-boundary value problems, and speculate on future developments.

Institution web page: http://www.unm.edu/

February 1, 2002

New Multivalue Methods for Differential Algebraic Equations

Minnie Kerr
North Carolina State University

Email: mykerr@unity.ncsu.edu

Abstract:

Multivalue methods are slightly different from the general linear methods John Butcher proposed over 30 years ago. Multivalue methods capable of solving differential algebraic equations (DAEs) have not been developed. In this paper, we have constructed three new multivalue methods for solving DAEs of index 1, 2, or 3, which include multistep methods and multistage methods as special cases. The concept of stiff accuracy will be introduced, and convergence results will be given based on the stage order of the methods. These new methods have the diagonal implicit property and thus are cheap to implement and will have order 2 or more for both the differential and algebraic components. We have implemented these methods with fixed step size, and they are shown to be very successful on a variety of problems. Some numerical experiments with these methods are presented.

Institution web page: http://www.ncsu.edu/

January 29, 2002

SiGMA: The Simulator Guided Memory Analyzer

Luiz DeRose
ACTC-IBM Research

Email: laderose@us.ibm.com

Abstract:

Application developers have been facing new and more complex performance tuning and optimization problems as parallel architectures become more complex. Users must manage the intricate interactions of many hardware and software components in order to achieve high performance. The efficient utilization of the memory subsystem is one of the critical factors to obtain performance on current and future architectures. However, the lack of performance tools that could help users to understand the behavior of the memory hierarchy makes it even more difficult for programmers to tune and optimize their codes.

In this talk I will present the Simulator Guided Memory Analyzer (SiGMA), an application performance tool under development at the Advanced Computing Technology Center, at IBM Research, for identification of problems, bottlenecks, and inefficiencies in a program caused by the mapping of data structures into the memory hierarchy. SiGMA is being designed to become an autonomic tool that will provide information on how well a code uses the memory system and guide application changes to improve the memory system performance. In addition, SiGMA is being designed to help performance prediction in current and future architectures.

Institution web page: http://www.research.ibm.com/actc/

January 14, 2002

Terascale Simulation of Trade Cumulus Convection

David Stevens

Lawrence Livermore National Laboratory

Email: dstevens@llnl.gov

Abstract:

This talk presents three-dimensional numerical simulations of oceanic trade cumulus clouds underlying stratocumulus clouds. The simulations are based upon a case studied in a Global Energy and Water Experiment Cloud System Study (GCSS) model intercomparison that is loosely based on observed conditions during the Atlantic Trade Cumulus Experiment (ATEX). It is motivated by the importance of this cloud type to global cloud radiative forcing, and their role as a feeder system for deep convection in the tropics. This study focuses on the sensitivity of the modeled cloud field to the domain size and the grid spacing. Domain widths from 6.5 to 20 kilometers and horizontal grid spacings ranging from 10 to 80 meters, with corresponding vertical grid spacing ranging from 5 to 40 meters, are studied, involving massively parallel computations on up to 2.5 billion grid cells. The combination of large domain size and small grid resolution provides an unprecedented perspective on this type of convection.

The mean stratocumulus cloud fraction, optical depth, and vertical fluxes of heat, moisture, and momentum are found to be quite sensitive to both the domain size and the resolution. The sensitivities are associated with a strong feedback between cloud fraction, cloud-top radiative cooling, and entrainment. The properties of individual cumulus clouds rising into the stratocumulus are less affected than the stratocumulus clouds. The simulations with 80 m horizontal/40 m vertical resolution are clearly underresolved, with distinctly different distributions of liquid water within the clouds. Increasing the resolution to finer than 40 m horizontal/20 m vertical affects the inversion structure and entrainment processes somewhat, but has less impact on the structure of individual clouds. Large-domain simulations exhibit mesoscale structure in the cloud organization and liquid water path. This mesoscale variability feeds back on the domain-mean properties through the cloud-radiative feedback. These simulations suggest that very large computations are required to obtain meaningful cloud statistics for this case.

Speaker's web page: http://www.llnl.gov/CASC/people/stevens/

Institution web page: http://www.llnl.gov/

December 20, 2001

An Overview of the DataFoundry Project

Terence Critchlow

Lawrence Livermore National Laboratory

Email: critchlow1@llnl.gov

Abstract:

The DataFoundry project is an ongoing effort to help scientists better interact with their data. As a result of this broad goal, the project has evolved over time to focus on different problems being faced by Lab scientists. This talk will outline, in an approximately chronological order, the various research areas we have been involved with. I will begin with an overview the project's original focus of data warehousing and integration, then I will briefly discuss each of the three efforts we are currently working on: large-scale data access, query pattern analysis, and ad hoc queries over mesh data.

Speaker's web page: http://www.llnl.gov/CASC/people/critchlow/

Institution web page: http://www.llnl.gov/

December 14, 2001

Introduction to the Sheaf Data Management System

David Butler
Limit Point Systems Inc.

Email: dmbutler@limitpt.com

Abstract:

The sheaf system is under development for the ASCI project to create both a theoretical data model and a concrete representation for scientific data that addresses the complexity of ASCI simulation data structures. The sheaf data model is based on mathematical sheaf theory, which incorporates topology, lattice theory, and graph theory. The sheaf data management system implements this model in C++, providing both a concrete representation for ASCI data and a wide range of powerful operations for manipulating data.

We begin by describing at a high level the fundamental functionality of the sheaf data model and architecture of the sheaf software system. We briefly review its history and development. We introduce a table + graph metaphor to describe the mathematical basis of the system. We then use the metaphor to discuss as many examples as time allows. Examples include conventional mesh-related applications, such as unstructured meshes, mesh queries, adaptive meshes, domain decomposition and parallelism, as well as less conventional, graph and order-related applications, such as inheritance hierarchies, version control, and project plans.

Institution web page: http://www.limitpt.com/index2.html

December 12, 2001

Variational Grid Adaptation Based on the Modified Equation Error Estimator

Giovanni Lapenta
Los Alamos National Laboratory

Email: lapenta@lanl.gov

Abstract:

A new method for grid adaptation will be presented. The method uses a variational approach based on a rigorous definition of the error.

In variational grid adaptation, appropriate functionals are minimized using the Euler–Lagrange approach. Typically the functionals are based on a heuristic understanding of the physics of the problem. A classic example is the Winslow approach where a merit function is used to adapt the grid to regions of strong gradients.

Recently, considerable attention has been devoted to using more mathematically rigorous definitions of the numerical error to construct automatic grid adaptation methods.

In the present seminar, I will show that a new approach can be proposed to combine the best and most reliable features of classic variational methods, such as the Winslow method, with more modern error estimators.

The seminar will be organized in a theoretical part and a series of examples.

The theoretical section will show how a rigorous definition of the error based on the modified equations approach can be used to derive new functionals for variational adaptation. The theoretical derivation will then be shown to lend itself to a simple practical implementation that can be used directly in pre-existing variational adaptation tools with minimal changes.

Examples will be shown to highlight the main features of the new approach. Elliptic problems and systems of Hyperbolic equations (gas dynamics and plasma physics) will be solved using adaptive grids. Particularly, I will compare the classic Winslow approach with the new method showing the remarkable improvements that can be obtained at no extra cost.

Institution web page: http://www.lanl.gov/

December 11, 2001

Multiscale Molecular Computation

Achi Brandt
Weizmann Institute of Science

Email: achi@wisdom.weizmann.ac.il

Abstract:

A general approach will be presented for describing macromolecules and fluids at increasingly coarser scales, and for deriving corresponding coarse-level Hamiltonians, based either on a given fine-scale Hamiltonian or, more fundamentally, on quantum chemistry. The coarse levels can be used for orders-of-magnitude acceleration of equlibrium (and also dynamics) simulations, avoiding slow-downs and attraction-basin traps typical to single-level simulations. They may further lead to "macroscopic equations," enabling simulations of indefinitely large systems.

Development of highly efficient multigrid methods for the level of electronic structure calculations will also be briefly surveyed.

Speaker's web page: http://www.wisdom.weizmann.ac.il/

Institution web page: http://www.weizmann.ac.il/

December 6, 2001

Subsetting and Clustering of array Comparative Genomic Hybridization Data

Annette Molinaro
UC Berkeley

Mark van der Laan, Dan H. Moore II

Email: molinaro@cc.ucsf.edu

Abstract:

Microarray-based comparative genomic hybridization is a unique way to measure genomic DNA aberrations and map them directly to the genome sequence. We will discuss the methodology behind this process and a statistical approach to the data. In this approach we will compare various ways of minimizing the number of variables by subsetting and discuss the results of clustering to predict tumor type.

Speaker's web page: http://oz.Berkeley.EDU/~laan/Students/Students_subpages/Annette/annette.html

Research web page: http://oz.Berkeley.EDU/~laan/Research/research.html

December 3, 2001

Enhanced Accuracy by Post-processing for Hyperbolic Problems

Bernardo Cockburn
University of Minnesota

Email: cockburn@math.umn.edu

Abstract:

We consider the enhancement of accuracy, by means of a simple post-processing technique, for finite element approximations to transient hyperbolic equations. The post-processing is a convolution with a kernel whose support has measure of order one in the case of arbitrary unstructured meshes; if the mesh is locally translation invariant, the support of the kernel is a cube whose edges are of size of the order of Delta x only. For example, when polynomials of degree k are used in the discontinuous Galerkin (DG) method, and the exact solution is globally smooth, the DG method is of order k+(1/2) in the L2 norm, whereas the post-processed approximation is of order 2k+1; if the exact solution is in L2 only, in which case no order of convergence is available for the DG method, the post-processed approximation converges with order k+(1/2) in L2(Omega_0) where Omega_0 is a subdomain over which the exact solution is smooth. Numerical results displaying the sharpness of the estimates are presented.

Speaker's web page: http://www.math.umn.edu/~cockburn/

Research web page: http://www.math.umn.edu/~cockburn/Research.html

Institution web page: http://www.umn.edu/

November 30, 2001

Using Hardware Monitors to Measure the Cache Eviction Behavior of Application Data Structures

Bryan Buck
University of Maryland

Email: buck@cs.umd.edu

Abstract:

This talk will discuss a proposed hardware performance monitor that provides support not only for measuring cache misses and the addresses associated with them, but also for discovering what data is evicted from the cache at the time of a miss. We will describe how to use this hardware support to efficiently measure the cache eviction behavior of application data structures at the source code level. We will then present the results of a study in which we implemented this technique in simulation and examined the overhead, perturbation of results, and usefulness of collecting this information. This talk will also discuss current developments on the Dyninst API project at the University of Maryland, which provides a library for modifying the code of running applications by inserting or removing instrumentation, and how it could be used in work like that presented here.

Speaker's web page: http://www.cs.umd.edu/users/buck/

Research web page: http://www.dyninst.org/

Institution web page: http://umd.edu/

November 29, 2001

A Streaming Supercomputer

William Dally and
Patrick Hanrahan
Stanford University

Email: billd@csl.stanford.edu

Abstract:

Modern VLSI technology enables 100GFLOPS chips (e.g., NVidea), 20 cent/MByte memory, and chips with 1Tbit/s of pin bandwidth, all of the ingredients for powerful and cost-effective scientific computing. Building supercomputers by clustering workstations and SMPs, however, doesn't realize the potential of the technology, instead giving machines that cost more per GFLOPS, GByte, and GUPS than their low-end counterparts.

Stream architectures have the potential to improve the performance (GFLOPS and GUPS) per unit cost by two orders of magnitude. Casting an application as a stream program makes all communication explicit, allowing much of it to be kept local (on chip), and hides the latency of global communication when it is needed. This in turn enables architectures with very high arithmetic intensity, e.g., 100GFLOPS chips with 20GBytes/s of memory bandwidth.

Realizing the performance and cost advantages of a streaming supercomputer, of course, requires recoding applications in a streaming style. We envision a layered programming system that simplifies the process of coding applications in this style and makes the resulting stream applications easily portable across many platforms.

hanrahan@Graphics.Stanford.edu

Speakers' web pages: http://csl.Stanford.EDU/~billd/

http://www-graphics.stanford.edu/~hanrahan/

Research web pages: http://www.graphics.stanford.edu

http://cits.stanford.edu

Institution web page: http://www.stanford.edu/

November 27, 2001

Immersive Virtual Reality for Scientific Visualization: A Progress Report

Andries van Dam

Thomas J. Watson, Jr. Professor,
Brown University

Email: avd@cs.brown.edu

Abstract:

Immersive virtual reality (IVR) has the potential to be a powerful tool for the visualization of burgeoning scientific datasets and models. While IVR has been available for well over a decade, its use in scientific visualization is relatively new, and many challenges remain before IVR can become a standard tool for the working scientist. In this presentation we provide a progress report and sketch a research agenda for the technology underlying IVR for scientific visualization. Among the interesting problem areas are how to do computational steering for exploration, how to use art-inspired visualization techniques for multi-valued data, and how to construct interaction techniques and metaphors for pleasant and efficient control of the environment. To illustrate our approaches to these issues, we will present specific examples of work from our lab, including immersive visualizations of arterial blood flow and medical imaging.

Speaker's web page: http://www.cs.brown.edu/people/avd/

Research web page: http://www.cs.brown.edu/research/

Institution web page: http://www.cs.brown.edu/

November 19, 2001

Temporal and Spatial Level of Details for Dynamic Meshes

Ariel Shamir
The Interdisciplinary Center

Email: arik@idc.ac.il

Abstract:

Multiresolution techniques enhance the abilities of graphics and visual systems to overcome limitations in time, space, and transmission costs. Numerous techniques have been presented that concentrate on creating level of detail models for static meshes. Time-dependent deformable meshes impose even greater difficulties on such systems. In this talk we describe a solution for using level of details for time-dependent meshes. Our solution allows for both temporal and spatial level of details to be combined in an efficient manner. By separating low- and high-frequency temporal information, we gain the ability to create very fast coarse updates in the temporal dimension, which can be adaptively refined for greater details.

Institution web page: http://www.idc.al.il/

November 16, 2001

Stochastic Multiresolution Models for Turbulence

Brandon Whitcher
National Center for Atmospheric Research

Email: whitcher@cgd.ucar.edu

Abstract:

The efficient and accurate representation of two-dimensional turbulent fields is of interest in the geosciences because the fundamental equations that describe turbulence are difficult to handle directly. Rather than extract the coherent portion of the image from the background variation, as in the classical signal plus noise model, we present a statistical model for individual vortices using the non-decimated discrete wavelet transform (MODWT). Unlike the orthonormal two-dimensional discrete wavelet transform (2D DWT), the 2D MODWT produces a redundant non-orthogonal transform. The main reason for this discrepancy is that the 2D MODWT does not downsample after convolving the wavelet filters with the input field. Hence, each wavelet field will have the same dimension as the original field.

A template image, supplied by the user, provides the features we want to extract from the observed field. By transforming the vortex template into the wavelet domain specific characteristics present in the template, such as size and symmetry, are broken down into components associated with spatial frequencies. Multivariate multiple linear regression is used to fit the vortex template to the observed vorticity field in the wavelet domain. Framing the regression model in the wavelet domain, through the multiresolution analysis, a template function that is idealized and simplistic in nature provides the foundation of a semi-parametric fit.

Application to a vortex census algorithm that tracks quantities of interest (such as size, peak amplitude, circulation, etc.) as the vorticity field evolves is given. Extensions to three dimensions are also possible.

Speaker's web page: http://www.cgd.ucar.edu/~whitcher/

Research web page: http://www.cgd.ucar.edu/stats/

Institution web page: http://www.cgd.ucar.edu/

November 16, 2001

Modelling Nuclear Contamination In Fractally Porous Media

Jim Douglas
Purdue University

Email: douglas@math.purdue.edu

Abstract:

The objects of this paper are to formulate a model for the transport of a chain of radioactive waste products in a fractured porous medium, to devise an effective and efficient numerical method for approximating the solution of the model, to implement the numerical method, and to present the results of some experimental calculations.

The derivation of the model begins with the single-porosity model considered in the thesis of Spagnuolo and first extends the model to a double-porosity model. Then, taking the time scale of the problem into account, this model is reduced to an analogue of the "limit model" for immiscible flow in a porous medium consdiered by Douglas, Paes, Leme, and Hensley. The well-posedness of the model follows from the theoretical results. A locally conservative Eulerian–Lagrangian numerical method is applied to the resulting system of equations; the experimental results cover both smooth and fractal distributions of permeability in the fractured reservoir.

Speaker's web page: http://www.math.purdue.edu/~douglas/

Institution web page: http://www.purdue.edu/

November 6, 2001

Morse Complexes and Topological Persistence

Herbert Edelsbrunner
Duke University

Email: edels@cs.duke.edu

Abstract:

We consider Morse complexes decomposing a manifold with a smooth height function into regions that have the same gradient flow pattern. We use a combinatorial algorithm with numerical components to construct such a complex via handle slides. A hierarchy of progressively simpler Morse complexes is then constructed by cancelling critical points in pairs. These cancellations are performed in the order of increasing importance, or persistence of critical points.

Speaker's web page: http://www.cs.duke.edu/~edels/

Institution web page: http://www.duke.edu/

October 29, 2001

Advances in Database Technologies for Emerging Applications

Walid Aref
Purdue University

Email: aref@cs.purdue.edu

Abstract:

The wide-spread use of the Internet and the explosion of modern applications demand new database engine technologies that accommodate the fluctuation, unpredictability,and new data types in data-intensive distributed information sources. The target is to develop and implement a broad class of innovative database technologies to address the demands of these emerging applications. In this talk, I will focus on data-intensive applications that handle distributed, spatial and multimedia information sources. I will present a collection of extensible database tools that address the needs of these applications. These include: high-dimensional space-filling curves and their applications in multimedia data scheduling, space-partitioning trees and their applications in indexing non-traditional data types in commercially used database systems, and pipelinable rank join algorithms for answering multi-feature nearest-neighbor queries that are common in multimedia database systems.

Speaker's web page: http://www.cs.purdue.edu/people/aref/

Research web page: http://www.cs.purdue.edu/homes/seh/PLACE.html

Institution web page: http://www.purdue.edu/

Coupled Eulerian-Lagrangian Simulations of Detonation Induced Shock Response in Tantalum

Daniel Meiron
California Institute of Technology

Email: dim@its.caltech.edu

Abstract:

The ASCI Alliance Center at Caltech is constructing a virtual shock physics facility with the aim of facilitating fully three dimensional simulations of the interaction of strong shock waves initiated by the detonation of high explosives (HE) with solid targets. In this talk we will present an overview of the current capabilities of the software environment under development at the center. A parallel implementation of an algorithm to dynamically couple Eulerian CFD (used for simulation of HE detonation) with Lagrangian solid mechanics will be presented along with computations on the ASCI terascale platforms utilizing this approach. We will also present results from a multi-scale modeling effort being pursued at the center to provide improved sub-grid scale models for those physical length and time scales that cannot be captured directly by simulation. We discuss the application of this approach to the development of a reduced reaction network for nitramine explosives and the development of models for plastic response of metals such as Tantalum.

Speaker's web page: http://www.acm.caltech.edu/people/faculty/meiron-d.html

Institution web page: http://www.caltech.edu/

October 24, 2001

Models for a Vertical Draining Film with an Insoluble Surfactant

Richard Braun
University of Delaware

Email: braun@math.udel.edu

Abstract:

The drainage of a thin Newtonian film with an insoluble surfactant is studied theoretically. Lubrication theory is applied to the thin film, which is suspended vertically from a frame and drains into a bath. In 1+1 dimensions, three nonlinear PDEs govern the evolution of the film shape, surface velocity and surfactant concentration. The surface viscosity and the surface tension of the films may be nonlinear functions of the surface concentration; we will consider linearized models primarily in this talk. Spatial discretization of the PDEs gives a large system of differential algebraic equations (DAEs) that may solved using perturbation methods, but results are typically obtained computationally. Slow and fast draining limits for the film drainage can be reached and a Marangoni-driven wave that is localized may be observed.

A model in 2+1 dimensions will also be presented that consists of 4 nonlinear PDEs and also requires the solution of a large set of DAEs. The results reveal an instability that is caused by a competition between gravity and Marangoni forces. The results obtained so far will be compared with experiment.

This work is in collaboration with S. Naire (now at U of Nottingham, UK) and S.A. Snow (Dow Corning) and has been supported by NSF DMS-9631287, DMS-9722854 and Dow Corning.

Speaker's web page: http://www.math.udel.edu/~braun/

Institution web page: http://www.udel.edu/

October 18, 2001

The Rate of Corrections and its Application in Scientific Computing

Zhiqiang Cai
Purdue University

Email: cai3@llnl.gov

Abstract:

One of the major issues in numerical analysis/scientific computing is the accuracy of the underlying numerical methods. It is a common knowledge that the accuracy of all current numerical methods for differentiation, integration, ordinary and partial differential equations, etc., are limited by the highest derivative of the underlying approximated function. Since solutions of many partial differential equations are not smooth, low order methods with adaptive mesh refinements seem to be the method of choice for such problems without the prior knowledge of singular behavior. But they are expensive. Obviously, it is desirable to develop higher order accurate methods for non-smooth problems. This seems to be a paradox.

This talk presents an approach for computing higher order accurate approximations for differentiation, integration, ODEs, and PDEs when the underlying approximated functions are not smooth. The key idea of this work is the introduction of the rate of corrections that is of universality, that quantifies the accuracy of the numerical method used, and that is computationally feasible. The rate of corrections is then used to increase the accuracy of the approximation. Also, this approach can be applied to problems without continuum background, such as the sequence of period-doubling bifurcations in chaos.

Speaker's web page: http://www.math.purdue.edu/~zcai/

Institution web page: http://www.purdue.edu/

October 17, 2001

Open Topics in Parallel I/O

Rob Ross
Argonne National Laboratory

Email: rross@mcs.anl.gov

Abstract:

Topics will include: implementing parallel i/o systems; design or choice of request languages; is atomicity always a requirement; and metadata issues in file system performance. I'm hoping that this talk will spark discussions during the rest of the week of my visit to LLNL.

Speaker's web page: http://www-unix.mcs.anl.gov/~rross/resume.htm

Research web page: http://www.parl.clemson.edu/pvfs/

Institution web page: http://www.anl.gov/

October 15, 2001

Mining Large Image Datasets

Jelena Tesic
UC Santa Barbara

Bangalore Manjunath

Email: jelena@ece.ucsb.edu

Abstract:

Mining large image datasets places a number of challenging requirements on the analysis framework. Some initial success has been achieved with systems that represent images as an organized collection of summarized information obtained from the feature descriptors and spatial constraints. However, the high dimensionality of the feature spaces and the size of the image datasets make meaningful summarization a challenging problem. A visual thesaurus based on low-level image descriptors provides a scalable conceptual framework for analyzing perceptual events. The heart of this method is a learning system that gathers information by interacting with database users.

Our main objective is to find clusters that represent similar feature points located in a small subset of a feature space. High-dimensional spaces represent challenges for clustering, due to the sparseness of the space. However, clusters may be formed from a couple of visually different elements that inhabit a large part of a high-dimensional space. Co-occurrence of clusters in an image helps us distinguish visually meaningful representatives. We are currently conducting experiments on texture feature sets to determine the dependency of the clusters in the texture feature space on feature vectors and spatial image layout. The objective of the visual thesaurus is to classify the image regions into perceptually similar categories. Spatial Event Cubes (SECs), are used represent and analyze the spatial relationships. SECs are computed with respect to particular spatial relationships. Detailed analysis shows that SECs can be used for visualization, discovery of latent spatial configurations, and for constructing indices for efficient and meaningful data access.

Speaker's web page: http://vision.ece.ucsb.edu/~jelena/

Research web page: http://www-iplab.ece.ucsb.edu/

Institution web page: http://www.ucsb.edu/

October 12, 2001

Optimization-Based Reference-Matrix Rezone Strategies for Arbitrary Lagrangian-Eulerian Methods on Unstructured Grids

Mikhail Shashkov
Los Alamos National Laboratory

Email: misha@t7.lanl.gov

Abstract:

The philosophy of the Arbitrary Lagrangian-Eulerian methodology (ALE; [1]) for solving multidimensional fluid flow problems is to move the computational grid, using the flow as a guide, to improve the accuracy and efficiency of the simulation. A principal element is the rezone phase in which a "rezoned" grid is created that is adapted to the fluid motion. Here we describe a general rezone strategy that ensures the continuing geometric quality of the computational grid, while keeping "rezoned" grid at each time step as close as possible to the Lagrangian grid. Although the methodology can be applied to more general grid types, we restrict ourselves here to unstructured triangular grids with fixed connectivity.

The unstructured grid is defined by positions of the nodes and connectivity. Between the node and each of its neighbors is an "edge-vector." Connectivity determines which edge-vectors form the columns of matrices (Jacobians) used to form the objective function minimized in the optimization.

The rezoning procedure consists of two phases: a sequence of local optimizations followed by a single global optimization. The local optimization leads to local "reference" Jacobians to be used in the global optimization. At each node we form a local patch from the Lagrangian grid and construct a local analog of a condition number functional [2]. Minimization of this functional with respect to position of central node defines its 'virtual' location (the node is not actually moved). By connecting this virtually-moved node to its (stationary) neighbors, we define "reference" edge-vectors and Jacobians that represent the best locally achievable geometric grid quality.

The "Rezoned" grid results from minimizing a global objective function, which measures the distance (in a least-squares sense) between the Jacobian of the "rezoned" grid and the "reference" Jacobian. This objective function includes a "barrier" that penalizes grids whose cells are close to being inverted. The global objective function is minimized by direct optimization. We provide numerical examples to demonstrate the robustness and the effectiveness of our methodology on model examples as well as for ALE calculations of hydrodynamics problems.

1. L. Margolin, Introduction to "Arbitrary Lagrangian-Eulerian Computing Method for All Flow Speeds", J. Comp. Phys., 135 (1997), pp. 198-202.

2. P. Knupp, Matrix Norms and the Condition Number - A General Framework to Improve Mesh Quality via Node-Movement, Proc. of the 8th Meshing Roundtable, S. Lake Tahoe CA, 1999, pp. 13-22.

Speaker's web page: http://cnls.lanl.gov/~shashkov

Institution web page: http://www.lanl.gov/

The ZSWEEP Algorithm for Rendering Irregular Grids

Claudio Silva
AT&T Labs-Research

Email: csilva@research.att.com

Abstract:

The need to visualize unstructured volumetric data arises in a broad spectrum of applications including structural dynamics, structural mechanics, thermodynamics, fluid mechanics, and shock physics. In this talk, I will focus on direct volume rendering, a term used to define a particular set of rendering techniques which avoids generating intermediary (surface) representations of the volume data. Instead, the scalar field is generally modeled as a cloud-like material, and rendered by computing a set of lighting equations.

I will focus on a particular technique called the ZSWEEP algorithm. The main idea of the ZSWEEP algorithm is very simple; it is based on sweeping the data with a plane parallel to the viewing plane in order of increasing depth (or "z"), "projecting" the faces of cells that are incident to vertices as they are encountered by the sweep plane. During ZSWEEP's face projection, we simply compute the intersection of the ray emanating from each pixel, and store their z-value, and other auxiliary information, in "sorted" order in a list of intersections for the given pixel, similar to the way an A-buffer operates. Compositing is performed as the "target Z" plane is reached. The efficiency arises from: (1) the fact that the algorithm exploits the implicit (approximate) global ordering that the z-ordering of the vertices induces on the cells that are incident on them, thus leading to only a very small number of ray intersection are done out of order; and (2) the use of early compositing which makes the memory footprint of the algorithm quite small.

In this talk, I will describe several features and (parallel, distributed, and out-of-core) extensions of the ZSWEEP algorithm. I will also discuss a technique for efficiently implementing a variant of ZSWEEP in hardware.

This work was done in collaboration with R. Farias (Mississippi State University), and J. Mitchell (State University of New York at Stony Brook).

Related papers can be obtained from the author's web page:

http://www.research.att.com/~csilva/papers

Speaker's web page: http://www.research.att.com/~csilva/

Institution web page: http://www.research.att.com/

October 11, 2001

Mimetic Finite Difference Methods for Diffusion Equations

Mikhail Shashkov
Los Alamos National laboratory

Email: misha@t7.lanl.gov

Abstract:

In this talk we review our ideas and present some new results related to construction of mimetic FDMs for solution of diffusion problems in strongly heterogeneous non-isotropic materials on nonorthogonal, nonsmooth, structured and unstructured computational grids. Development of high-quality finite-difference methods (FDMs) for diffusion equation is a part of a bigger effort to create a discrete analog of vector and tensor calculus, that can be used to accurately approximate continuum models for a wide range of physical processes. These FDMs preserve fundamental properties of the original continuum differential operators and allow the discrete approximations of partial differential equations (PDEs) to mimic critical properties including conservation laws and symmetries in the solution of the underlying physical problem.The discrete analogs of differential operators satisfy the identities and theorems of vector and tensor calculus and provide new reliable algorithms for a wide class of PDEs.

We describe discrete analogs of divergence and flux operators and discuss their properties. First, we describe basic (global) FDM, which uses these operators. Basic method uses cell-centered discretization for scalar unknowun and normal component of flux on each face. Next, we describe local modification of our method where we have addtional unknowns on the faces of the cell. Then, we present possible generalizations,which includes more general boundary conditions, unstructured grids, case of three dimensions, cylindrical coordinates and so on. Finally, we present some numerical examples.

Speaker's web page: http://cnls.lanl.gov/~shashkov

Institution web page: http://www.lanl.gov/

October 3, 2001

Using Visualization to Understand the Behavior of Computer Systems

Robert Bosch Jr.
Stanford University

Email: bosch@cs.stanford.edu

Abstract:

This talk introduces Rivet, a general-purpose environment for the development of computer systems visualizations. Rivet can be used for both real-time and post-mortem analyses of data from a wide variety of sources. The modular architecture of Rivet enables sophisticated visualizations to be assembled using simple building blocks representing the data, the visual representations, and the mappings between them. The implementation of Rivet enables the rapid prototyping of visualizations through a scripting language interface while still providing high-performance graphics and data management.

Rivet has been used for a wide variety of studies, ranging from processors and memory systems to multiprocessors and clusters to wide-area networks. The talk includes two examples. The first, a visualization displaying the behavior of superscalar processor pipelines, demonstrates the value of coordinated multi-view visualizations for understanding computer systems components. The second, a detailed performance analysis of the Argus parallel graphics library, shows how Rivet can be used in conjunction with comprehensive data sources such as the SimOS complete machine simulator to provide a powerful iterative analysis framework for understanding computer systems as a whole.

Speaker's web page: http://www-flash.stanford.edu/~bosch/

Research web page: http://graphics.stanford.edu/projects/rivet/

Institution web page: http://www.stanford.edu/

October 1, 2001

Application of Probe-Based Storage to High Performance Computing

Zachary Petersen
UC Santa Cruz

Darrell Long

Email: zachary@cse.ucsc.edu

Abstract:

Although there is continual improvement in speed and capacity of storage systems, traditional longitudinal magnetic recording is approaching a hard physical limit. A performance gap between the speed and capacity of RAM and disk is increasing at a rate of 50% a year. This performance gap is especially visible in high-performance workloads such as those employed in NNSA's Advanced Simulation and Computing Program (ASCI). We in the Computer Systems Research Group at the University of California, Santa Cruz are investigating two lines of new and innovative high-performance storage research.

The first involves an exciting new storage technology based on MicroElectroMechanical Systems (MEMS) which promises a significant increase in performance, capacity, and reliability relative to modern storage devices. The second, approaches the performance problems of current storage devices, applying novel ideas of grouping like data to improve read performance, while eliminating the amount of rewriting needed to be done to keep these data contiguous. Using sanitized workload traces from LLNL and support from the ISCR program, we have shown both lines of research successful in addressing the I/O problems posed by ASCI applications.

Speaker's web page: http://www.cse.ucsc.edu/~zachary/

Research web page: http://csl.cse.ucsc.edu/

Institution web page: http://www.ucsc.edu/public/